Multi-Layer Anatomically-Guided Diffusion and Adaptive Calibration for Precision Radiotherapy Dose Prediction

Lei Xu, PhD; Kunpeng Zhang, MS; Jinlu Ma, MD

Department of Radiology, The First Affiliated Hospital of Xi'an Jiaotong University, Xi'an, Shaanxi, 710061, China

2026 JULY 19–22 VANCOUVER, BC

JOINT AAPM | COMP MEETING

ABSTRACT

Purpose:
Dose prediction can improve radiotherapy treatment planning efficiency, but existing methods often insufficiently model inter-slice anatomical information and show limited dose accuracy within clinically critical regions of interest (ROIs). This study proposes MAGDiff-ADC, a diffusion-based framework integrating multi-layer anatomical guidance and adaptive ROI-focused calibration for accurate dose prediction.

Methods:
MAGDiff-ADC formulates dose prediction as a conditional denoising diffusion process, progressively generating dose maps from noise under the guidance of CT images and ROI masks. A conditional U-Net with multi-layer spatially adaptive normalization hierarchically incorporates adjacent-slice anatomical features to capture local 3D dose–anatomy correlations. A lightweight fully convolutional Adaptive Dose Calibration module further refines ROI dose values to reduce pixel-level inaccuracies. The framework was evaluated on an in-house dataset and a public benchmark using multiple dosimetric and image-level metrics.

Results:
MAGDiff-ADC achieved state-of-the-art performance across both datasets. Multi-layer anatomical guidance improved inter-slice consistency and anatomical fidelity, while the ADC module enhanced dose accuracy within clinically important ROIs.

Conclusion:
MAGDiff-ADC provides an effective diffusion-based approach for radiotherapy dose prediction by combining adjacent-slice anatomical guidance with ROI-targeted calibration, improving both global dose distribution quality and ROI-specific dosimetric accuracy.

INTRODUCTION

Radiotherapy treatment planning relies on an optimized three-dimensional dose distribution that maximizes planning target volume (PTV) coverage while sparing organs at risk (OARs). In routine practice, plan optimization is still performed manually in a trial-and-error manner, which is time-consuming, planner-dependent, and variable across institutions. Automated voxel-wise dose prediction can provide a high-quality reference dose map to accelerate optimization, standardize plan quality, and enable rapid OAR risk assessment.

Recent knowledge-based planning and deep learning approaches have improved prediction performance, but two issues still limit clinical utility. First, many methods process dose prediction slice-by-slice and do not explicitly exploit inter-slice anatomy that governs volumetric dose continuity. Second, even when global similarity is high, pixel-level dose errors within clinically critical regions of interest (ROIs), especially near steep PTV–OAR interfaces, may remain non-negligible.

To address these limitations, we propose MAGDiff-ADC, a two-stage framework that combines a context-aware conditional diffusion model for anatomically faithful dose generation with a lightweight Adaptive Dose Calibration (ADC) module for ROI-focused numerical refinement.

METHODS AND MATERIALS

MAGDiff-ADC consists of a multi-layer anatomically guided diffusion network and an ROI-aware adaptive dose calibration module (Fig. 1). MAGDiff formulates radiotherapy dose prediction as a conditional denoising diffusion process. During training, reference dose maps are progressively corrupted with Gaussian noise, and a conditional U-Net learns to recover the clean dose by predicting the injected noise. CT images and ROI masks, including PTVs and OARs, are used as anatomical conditions. The U-Net backbone incorporates time-embedded residual blocks and self-attention to capture diffusion-step information and long-range spatial dependencies.

To incorporate local 3D anatomical context while avoiding the computational burden of full 3D diffusion, we introduce Multi-Layer Spatially-Adaptive Normalization (MLSAN). MLSAN learns spatially varying affine parameters from the current slice and its adjacent slices, and injects multi-slice CT and ROI information into decoder features at multiple scales. This design enables the model to better preserve inter-slice dose continuity and dose–anatomy consistency.

Because stochastic diffusion sampling may still introduce local numerical deviations, an adaptive dose calibration module is further used for deterministic refinement. This lightweight ResUNet takes the preliminary diffusion-predicted dose and anatomical conditions as inputs, refines dose values within clinically relevant ROIs, and blends the calibrated ROI dose back into the full dose map using a softened ROI mask. MAGDiff and ADC are trained sequentially using denoising loss and ROI dose reconstruction loss, respectively.

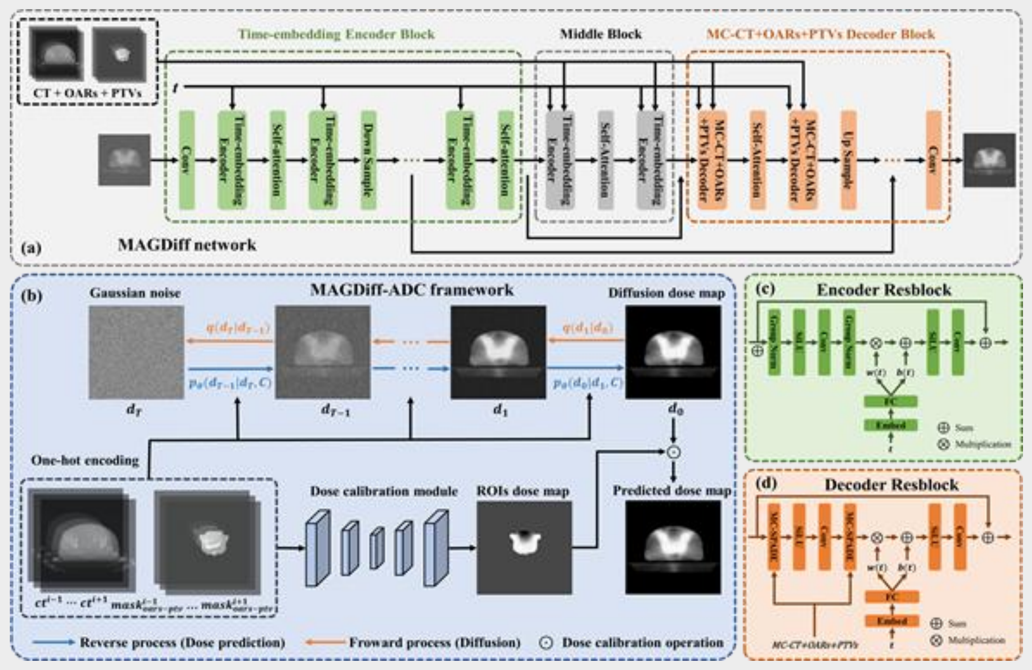

Figure 1. Overview of MAGDiff-ADC. MAGDiff performs conditional diffusion-based dose generation with MLSAN-based multi-slice anatomical modulation, followed by ADC for ROI-focused dose calibration.

EXPERIMENTS

We evaluated MAGDiff-ADC on two datasets. (1) An in-house cervical cancer IMRT dataset comprising 110 patients (prescription 50 Gy) with CT, PTV, and OAR masks (rectum, bladder, spinal cord, bilateral femoral heads, and bilateral kidneys); 74 cases were used for training and 36 for retrospective testing. (2) The public AAPM OpenKBP head-and-neck dataset including CT, multi-level PTV prescriptions, and standard OARs.

Performance was assessed using image-level metrics (mean absolute error, MAE; structural similarity, SSIM) and clinically relevant criteria, including the Dose Score and DVH Score (absolute differences in DVH criteria).

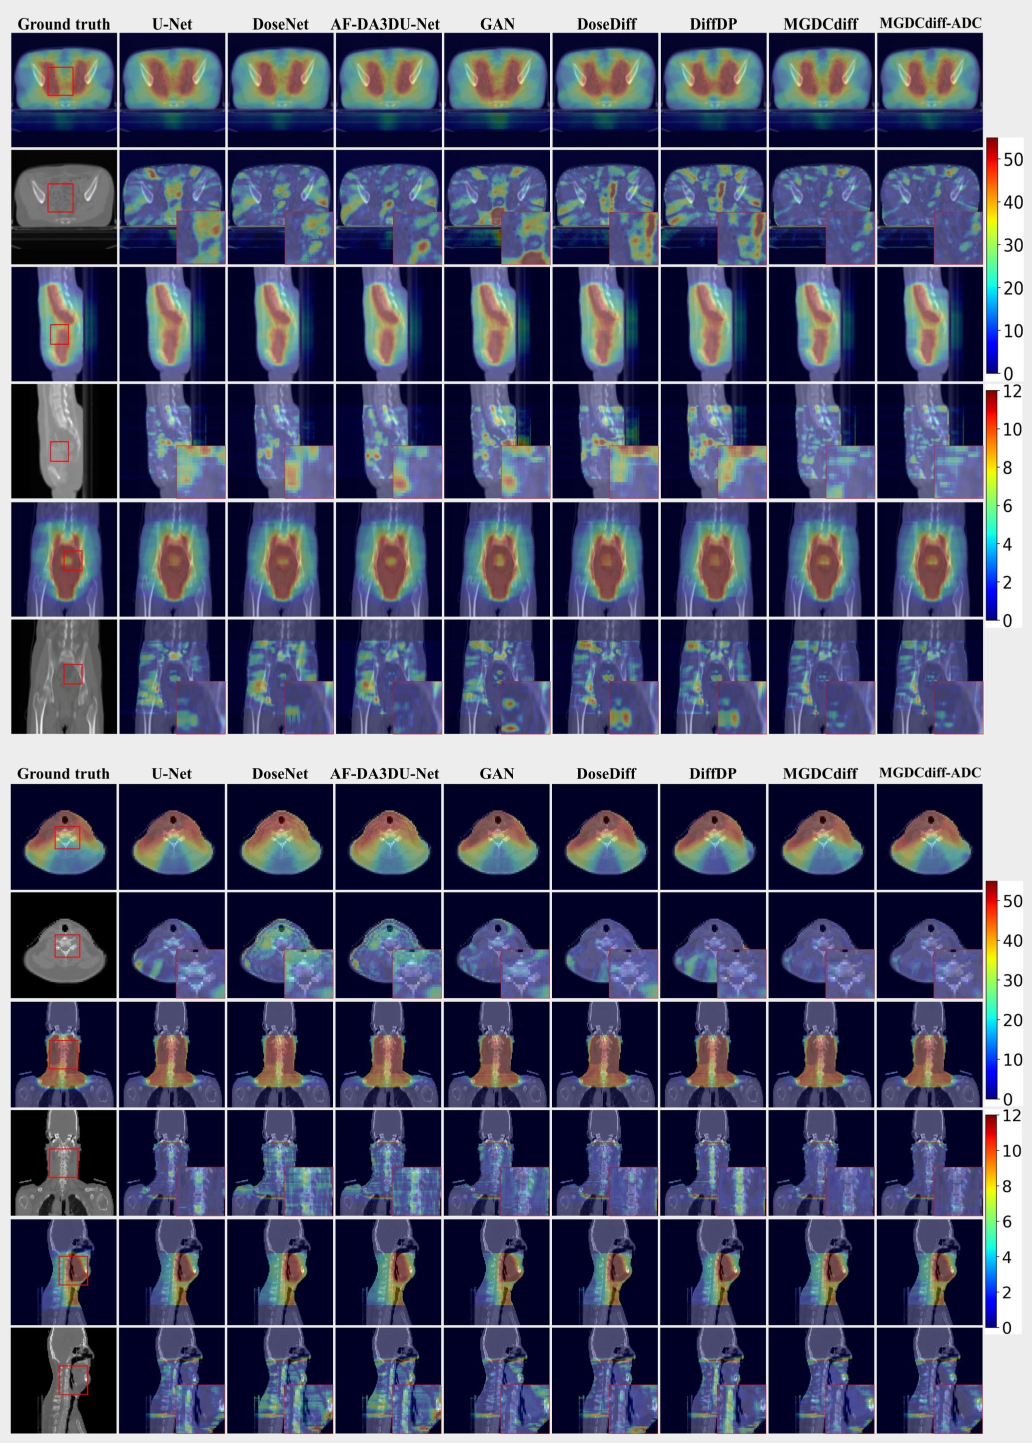

Figure 2. The visual comparison of predicted dose distribution maps obtained with different methods on the cervical cancer dataset and OpenKBP data.

RESULTS

Across both datasets, MAGDiff-ADC achieved state-of-the-art performance, improving overall dose fidelity and DVH agreement. On the cervical cancer test set, it produced dose maps visually closest to the reference (Fig. 2) and achieved Dose Score 2.211±0.130 Gy, DVH Score 0.734±1.105, and SSIM 0.852±0.015. On OpenKBP, it also outperformed competing methods with Dose Score 4.173±1.065 Gy, DVH Score 1.259±1.921, and SSIM 0.924±0.038.

Ablation analyses demonstrated complementary contributions from MLSAN and ADC (Fig. 3). Replacing naive feature concatenation with spatially adaptive modulation improved prediction quality, and injecting adjacent-slice anatomy further reduced Dose and DVH errors, confirming the importance of explicit inter-slice guidance. Finally, ADC consistently improved ROI-level accuracy and DVH alignment by correcting localized deviations from diffusion sampling, particularly in sensitive OARs near high-gradient regions.

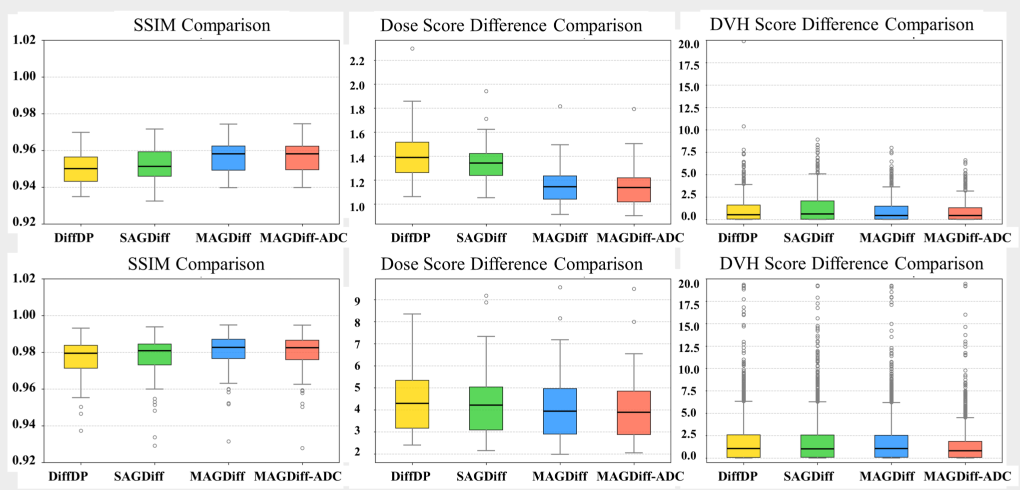

Figure 3. Ablation study on cervical cancer (top) and OpenKBP (bottom). Boxplots summarize SSIM and Dose/DVH score differences for DiffDP, SAGDiff, MAGDiff, and MAGDiff-ADC.

CONCLUSIONS

We presented MAGDiff-ADC, a multi-layer anatomically guided diffusion framework with adaptive ROI calibration for radiotherapy dose prediction. By combining MLSAN-based inter-slice anatomical fusion with an efficient calibration network, the proposed method improves both global dose fidelity and clinically critical ROI accuracy on an in-house cervical cancer cohort and the OpenKBP benchmark, strengthening the clinical usefulness of automated dose prediction for efficient and consistent treatment planning.

REFERENCES

1. Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. Advances in neural information processing systems, 33, 6840-6851.
2. Feng, Z., Wen, L., Wang, P., Yan, B., Wu, X., Zhou, J., & Wang, Y. (2023, October). DiffDP: Radiotherapy dose prediction via a diffusion model. In International conference on medical image computing and computer-assisted intervention (pp. 191-201). Cham: Springer Nature Switzerland.
3. Zhang, Y., Li, C., Zhong, L., Chen, Z., Yang, W., & Wang, X. (2024). DoseDiff: distance-aware diffusion model for dose prediction in radiotherapy. IEEE Transactions on Medical Imaging, 43(10), 3621-3633.
4. Babier, A., Zhang, B., Mahmood, R., Moore, K. L., Purdie, T. G., McNiven, A. L., & Chan, T. C. (2021). OpenKBP: the open-access knowledge-based planning grand challenge and dataset. Medical Physics, 48(9), 5549-5561.

CONTACT

Lei Xu, PhD
Department of Radiology, The First Affiliated Hospital of Xi'an Jiaotong University
Address: No.277 Yanta West Road, Xi'an, Shaanxi, 710061, China
radiumxu@xjtufh.edu.cn
+86 17706446660